AI Based Automated Traffic Monitoring System for Vehicles and License Plate Recognition

***Read DANJOLLI*[1]**

Abstract

This paper presents a comprehensive architecture for automated traffic violation surveillance. It is based on sophisticated deep learning algorithms and artificial intelligence systems with computer vision. The main objective is to develop an integrated pipeline that integrates vehicle detection, Automatic License Plate Recognition (ALPR), and visual attribute classification (e.g., color, manufacturer, and model). YOLO detection, DeepSORT tracking, CRNN network OCR, and CNN for car brand and color categorization are all parts of the technical solution. The study fully compares Edge and Cloud architectures, examining how well they perform under different conditions, such as high traffic and poor lighting. The findings show that, while Cloud solutions offer more flexibility but at a higher latency cost, Edge solutions, despite their processing limitations, achieve response times below 200 ms and accuracy above 95% in license plate identification.

Along with specific implementation recommendations for the Albanian context, the study addresses algorithmic fairness, privacy protection and GDPR compliance. It also addresses the ethical and legal elements of using surveillance technologies, highlighting the prospects and challenges for a successful adoption in Albania.

Furthermore, to compensate for the personalized data pages for the Albanian market, synthetic data models were included in the initial training. This was

[1] Read Danjolli graduated in Information Technology – Economic Informatics at the European University of Tirana. He also holds a master's degree from the Faculty of Engineering, Informatics and Architecture at the European University of Tirana, in Computer Engineering (MSc), Software Engineering profile.

sufficient on the ground to allow for higher algorithmic adaptability. Investments in human resource training and a well-defined framework are also necessary for the deployment of technologies to ensure accountability, transparency and responsibility for all. This comprehensive strategy lays the foundation for an automated application system in Albania that is reliable and sustainable.

Key words: *Automated Traffic Surveillance, Automatic License Plate Recognition, Vehicle Attribute Recognition, Deep Learning & Computer Vision, Edge and Cloud Architecture, Ethical & GDPR Compliance*

Introduction

Introduction and Background

Road safety and efficient traffic management are fundamental pillars of modern societies. The ability of a nation to ensure safe mobility directly impacts its social welfare, economic productivity, and environmental sustainability (World Health Organization, 2023). Albania, like many developing economies, has experienced a dramatic increase in the number of registered vehicles in the past three decades. More than 700,000 vehicles are now circulating in a relatively small and urbanizing country, reflecting greater economic access to transportation but also introducing significant costs in terms of congestion, urban pollution, and rising accident rates (European Transport Safety Council, 2022).

Statistics show that thousands of traffic incidents are recorded annually in Albania, with dozens of fatalities and hundreds of serious injuries. These events translate into economic losses exceeding 2.5% of the national GDP, through healthcare costs, productivity losses, infrastructure damage, and delays in mobility (INSTAT, 2023). Beyond their economic weight, these incidents expose the inability of current institutional frameworks to regulate and enforce traffic laws effectively. Police patrols and conventional monitoring systems lack the scalability and precision to handle the complexities of modern traffic in Albania's rapidly growing cities.

This situation places Albania in a unique position. On one hand, it suffers from the absence of a centralized and automated system for identifying and sanctioning traffic violations. On the other hand, its manageable size and urgent need for digitalization create an opportunity to adopt cutting-edge intelligent surveillance technologies. By integrating Artificial Intelligence (AI), Computer Vision, and Automated License Plate Recognition (ALPR) systems, Albania could leapfrog traditional approaches and establish a transparent, efficient, and accountable traffic management ecosystem.

However, such a transformation cannot occur in a vacuum. Ethical, legal, and social concerns must accompany technological progress. AI-driven surveillance brings questions of privacy, fairness, and compliance with European regulations such as the General Data Protection Regulation (GDPR). Ensuring public trust, institutional accountability, and legal legitimacy will be critical for success (European Union Agency for Fundamental Rights, 2021).

Problem Statement

The core research problem addressed in this study lies in designing an integrated pipeline capable of recognizing vehicle license plates and associated attributes—such as color, make, and modeling real-world conditions. This problem is multidimensional, cutting across technical, infrastructural, and ethical domains.

ALPR and Vehicle Attribute Recognition (VAMR) systems must function reliably in diverse scenarios: poor lighting, nighttime conditions, adverse weather (rain, fog, snow), high vehicle speeds, occlusions and damaged or non-standard plates. Such environmental and operational variables significantly challenge AI algorithms, which often underperform when exposed to conditions outside their training datasets (Redmon & Farhadi, 2018).

Albania faces an additional complexity due to the coexistence of multiple license plate formats, national, historical, and foreign. Flexible algorithms are required to recognize this heterogeneity, including multilingual fonts, varying colors, and design elements. Similarly, vehicles on Albanian roads exhibit great variety in brand and model, often making fine-grained classification difficult. Most existing solutions treat license plate recognition and vehicle attribute recognition separately. The challenge is to integrate these tasks into a single coherent architecture, where the outputs of one module reinforce the reliability of the other. This raises architectural and algorithmic issues, particularly concerning error propagation and confidence fusion between subsystems (Li et al., 2019).

Law enforcement applications demand real-time inference, often across multiple cameras deployed simultaneously. This introduces trade-offs between computational cost, latency, and system scalability. Choosing between Edge computing, where data is processed locally on embedded devices, and Cloud computing, where data is centralized for processing, is a major design decision.

High-performing Deep Learning models require large, high-quality annotated datasets. Albania lacks curated traffic surveillance datasets tailored to its specific conditions. Manual data collection and annotation is resource-intensive, while synthetic data generation and transfer learning only partially mitigate this gap.

Finally, AI-based traffic surveillance inevitably raises ethical dilemmas. License plates are legally considered personal data under the GDPR, and improper storage or processing risks violating individual rights. Biases in algorithmic performance could disproportionately impact minority groups, while misidentifications could lead to unjust sanctions (Barocas, Hardt, & Narayanan, 2019). Addressing these concerns requires robust anonymization techniques, transparent auditing mechanisms, and public communication strategies.

Objectives, research questions and hypothesis

In response to these challenges, this study sets fifth the following objectives:

- To analyze and evaluate the application of state-of-the-art AI techniques, especially Deep Learning, in ALPR and VAMR tasks under Albanian traffic conditions.
- To design an integrated architecture that combines license plate recognition and vehicle attribute recognition in a unified pipeline.
- To develop or adapt Deep Learning models capable of identifying vehicle color, make, and model with high accuracy.
- To assess system performance using standardized metrics, focusing on accuracy, speed (frames per second), and robustness in real-world conditions.
- To examine ethical, legal, and social implications of deploying automated surveillance in Albania, ensuring compliance with GDPR and public trust.

Research Questions

Building on the objectives, the study frames its inquiry around several guiding research questions (RQs):

RQ1: What are the most effective Deep Learning architectures for implementing ALPR and VAMR in high-volume traffic conditions?
RQ2: How can the robustness of ALPR systems be improved to handle low lighting and adverse weather conditions?
RQ3: What is the optimal architecture for integrating ALPR and VAMR modules into a unified pipeline, balancing accuracy and computational efficiency?
RQ4: What are the trade-offs between Edge and Cloud inference for real-time automated traffic surveillance?
RQ5: What ethical, legal, and privacy considerations must be addressed for large-scale deployment in Albania?

The main **hypothesis** guiding this research is: ***AI-driven solutions significantly enhance speed, accuracy, and reliability in traffic violation detection compared to traditional systems, thereby improving road safety and enforcement efficiency.***

Literature Review and Gaps

Evolution of Traffic Surveillance Technologies

Traffic monitoring has undergone a remarkable transformation over the past century. Early methods were entirely manual, with police officers stationed at intersections to observe violations and direct traffic flow. In the mid-20th century, pneumatic tubes and inductive loop detectors (ILDs) represented the first steps toward automation, offering basic vehicle counts and speed measurements.

The proliferation of Closed-Circuit Television (CCTV) in the 1960s–1970s introduced visual monitoring, but systems still relied heavily on human operators. In the 1980s–1990s, Video Incident Detection (VID) systems emerged, using background subtraction and motion analysis to detect sudden stops, wrong-way driving, and congestion (Parker & Harris, 1998). Yet these were fragile under lighting variation and weather changes.

From the 1990s onwards, radar and LiDAR became widely used for speed enforcement and red-light violation detection, enabling partial automation. These systems, however, were limited in their ability to identify vehicles beyond plate numbers, offering little in terms of broader traffic analytics (Zhang et al., 2017).

The 2010s marked the era of **AI-powered traffic surveillance,** where Convolutional Neural Networks (CNNs) revolutionized object detection and recognition. Today, AI-based ALPR/VAMR systems can detect vehicles, read plates, and classify attributes (color, make, model) with high accuracy in near real-time, representing a leap in automation, scalability, and analytical depth (Redmon & Farhadi, 2018).

Artificial Intelligence in Intelligent Transportation Systems (ITS)

AI integration has significantly expanded the capabilities of Intelligent Transportation Systems (ITS). Its strengths lie in processing vast amounts of heterogeneous data—from cameras, sensors, GPS, and mobile applications—to support predictive modeling, real-time decision-making, and autonomous operation. Applications include:

- Traffic flow prediction: Machine Learning models forecast congestion using time-series analysis
- Adaptive traffic light control: Reinforcement Learning optimizes signal timings in real time.
- Navigation systems: Platforms like Google Maps and Waze apply AI to predict travel times and reroute drivers dynamically.

- Autonomous driving: Deep Learning enables environment perception, sensor fusion, and decision-making in self-driving cars.
- Law enforcement: Computer Vision detects violations such as speeding, red-light running, illegal parking, and distracted driving (Li et al., 2019).

AI thus transforms ITS from reactive infrastructures into proactive, intelligent ecosystems that optimize safety, efficiency, and sustainability.

FIGURE 1: Example of vehicle and license plate recognition.

Automated License Plate Recognition (ALPR)

ALPR is a cornerstone technology in traffic enforcement. It typically involves four stages: image acquisition, plate detection, character segmentation, and optical character recognition (OCR). Advances in object detection, particularly YOLO (You Only Look Once), SSD (Single Shot Detector), and Faster R-CNN, have greatly improved plate localization. OCR accuracy has been boosted by CRNNs (Convolutional Recurrent Neural Networks), which combine CNNs and LSTMs for end-to-end recognition without explicit character segmentation (Shi et al., 2017).

Challenges persist, however, in dealing with varying plate designs, occlusions, glare, damaged characters, and high-speed motion blur. Benchmarks show state-of-the-art ALPR systems achieving 95–98% accuracy under controlled conditions but lower performance in adverse real-world scenarios.

Vehicle Attribute Recognition (VAMR)

While license plates uniquely identify vehicles, attribute recognition enhances reliability and utility. Color, make, and model recognition aid in detecting cloned or stolen plates, re-identifying vehicles across camera views, and enabling traffic composition analysis (Sochor et al., 2018).

Fine-grained classification is the key challenge: distinguishing between visually similar models or color shades requires high-quality datasets and sophisticated CNN architectures. Large,47784 annotated datasets such as CompCars and BoxCars116k have accelerated progress (Yang et al., 2015). Transfer learning from models trained on ImageNet has also proven highly effective.

Ethical and Privacy Concerns

The expansion of surveillance technologies has sparked extensive debate on privacy and civil liberties. ALPR inherently creates location trails of vehicles, which, if stored extensively, risk enabling mass surveillance beyond traffic law enforcement. GDPR explicitly classifies license plates as personal data, requiring strict adherence to principles of data minimization, storage limitation, and purpose specification.

Algorithmic fairness is another critical concern. Biases in training data can produce disproportionate misidentification rates across demographic groups, undermining trust and legitimacy (Buolamwini & Gebru, 2018). Transparent auditing, explainability, and public engagement are increasingly emphasized as safeguards (European Commission, 2021).

Despite notable advancements, several research gaps persist. First, there is a lack of integrated evaluations, as most studies analyze ALPR and VAMR systems separately rather than as unified, real-world pipelines. Second, regional adaptation remains limited, with datasets primarily reflecting Western formats and failing to generalize to regions like Albania. Third, edge–cloud trade-offs are insufficiently explored, particularly regarding latency, cost, and privacy balance. Fourth, the field suffers from benchmarking inconsistencies, as new Deep Learning architectures emerge faster than standardized evaluations. Finally, ethical considerations are often discussed broadly, without adaptation to specific national and cultural contexts.

Computer Vision and Deep Learning Foundations

Computer Vision (CV) provides the theoretical and practical foundation for automated traffic surveillance. Its objective is to enable machines to interpret visual information from images and video streams. Historically, CV relied on handcrafted feature extraction methods such as Scale-Invariant Feature Transform (SIFT) and Histogram of Oriented Gradients (HOG), which required human expertise to design features relevant for object detection.

The emergence of **Deep Learning**, particularly Convolutional Neural Networks (CNNs), revolutionized the field by enabling end-to-end learning. CNNs automatically extract hierarchical visual features, from low-level edges to high-level object representations, thereby surpassing handcrafted approaches in robustness and scalability (LeCun, Bengio, & Hinton, 2015). For sequence-based tasks, Recurrent Neural Networks (RNNs) and their variants, such as Long Short-Term Memory (LSTM), handle temporal dependencies in traffic video streams. Together, CNN-RNN hybrids such as CRNNs are particularly effective in character recognition for license plates (Shi et al., 2016).

Data Requirements for Surveillance Models

Deep learning systems rely heavily on large and well-curated datasets to achieve effective training and reliable performance. The size, diversity, and quality of the dataset play a crucial role in determining how well a model can generalize to new, unseen scenarios. To ensure this generalization, several factors must be carefully considered during dataset design and preparation.

One of the most important aspects is the diversity of environments represented in the dataset. Models trained only under ideal conditions often fail when faced with real-world variability. Therefore, datasets should include images captured in different lighting situations—day and night—as well as under various weather conditions such as rain, fog, or snow. Additionally, the inclusion of different traffic densities and urban or rural scenes enhances the system's ability to perform consistently across diverse environments.

Another key consideration is balance representation. In traffic surveillance and vehicle recognition systems, certain vehicle types, such as sedans or compact cars, are much more common than others, like buses, motorcycles, or trucks. If a dataset reflects this imbalance, the model may become biased, performing well in frequent classes but poorly on rare ones. Ensuring an even distribution of vehicle categories helps maintain fairness and accuracy across all types.

To further strengthen dataset quality, data augmentation techniques are widely applied. Methods such as random cropping, image rotation, flipping, and brightness adjustments can artificially increase the size of the dataset. These transformations expose the model to a broader range of visual variations, making it more robust to distortions, angles, and lighting changes that occur in real-life traffic footage.

In addition, transfer learning provides an efficient solution for overcoming data limitations. Instead of training a model entirely from scratch, researchers can fine-tune neural networks that have already been pretrained on large-scale datasets like ImageNet. This approach allows models to benefit from previously learned visual features and significantly reduces the need for massive amounts of new, domain-specific data (Kornblith et al., 2019).

Another promising technique is the use of synthetic data computer-generated imagery that replicates real-world traffic scenes. Synthetic datasets can simulate different conditions, vehicle types, and perspectives that may be difficult or expensive to capture manually. However, successful integration of synthetic and real data requires careful domain adaptation, ensuring that models trained on artificial images can perform effectively when applied to real environments.

In the context of Albania, one of the major challenges in developing deep learning-based traffic surveillance systems is the lack of annotated, high-quality datasets. To address this, a hybrid strategy is recommended. This involves combining synthetic datasets with transfer learning approaches while gradually collecting and annotating real-world images from Albanian traffic environments. Such a method not only accelerates initial system development but also lays the foundation for continuous improvement as more local data becomes available.

Through this integrated approach—balancing diversity, augmentation, transfer learning, and synthetic generation—deep learning models for traffic surveillance in Albania can achieve greater accuracy, adaptability, and long-term scalability.

Preprocessing Pipelines: Detection and Tracking

Traffic surveillance systems operate through structured pipelines that begin with the detection of vehicles and license plates and proceed with tracking them across consecutive frames. This process ensures that every detected object is not only recognized but also consistently followed over time, allowing accurate monitoring of traffic flow and potential violations.

The first step, object detection, focuses on identifying vehicles or license plates within individual frames. One of the most influential models in this domain is YOLO (You Only Look Once), which treats detection as a single regression problem over bounding boxes and class probabilities (Redmon & Farhadi, 2018). YOLO's remarkable speed and efficiency make it ideal for real-time traffic analysis, where high frame rates and immediate detection are essential.

Following detection, object tracking maintains the continuity of identified vehicles across frames. A leading approach is DeepSORT, which enhances the original SORT algorithm by combining Kalman filtering with deep feature embeddings (Wojke, Bewley, & Paulus, 2017). This enables the system to preserve the identity of vehicles even when they temporarily disappear due to occlusions or overlaps.

Together, these methods form a "tracking-by-detection" architecture, an integrated framework in which detection provides bounding boxes for each object, and tracking ensures their consistent identification throughout the video sequence. This synergy between detection and tracking is fundamental for building reliable, real-time traffic surveillance systems capable of continuous and accurate vehicle monitoring.

Key Algorithms and Architectures

TABLE 1: Summary of key algorithms used in ALPR and vehicle attribute recognition systems.

Algorithm	Architecture Type	Primary Role	Strengths	Challenges
YOLO (v3–v5)	CNN, one-stage detector	Vehicle and plate detection	High speed, real-time performance	Struggles with small or distant objects
DeepSORT	Kalman filter + CNN embeddings	Multi-object tracking	Robust to short-term occlusion, preserves identity	Requires high-quality detectors, failures under long occlusion
CRNN	CNN + RNN/ LSTM hybrid	License plate OCR	End-to-end recognition without segmentation	Sensitive to low-quality plate images
ResNet, VGG	CNN classifiers	Vehicle attribute recognition (color, make, model)	Strong classification ability, transfer learning	Requires large, annotated datasets
CTC Loss	Loss function	Sequence prediction training	Enables flexible sequence alignment	Requires careful hyperparameter tuning

Together, these algorithms form the backbone of an integrated ALPR+VAMR system, capable of high-accuracy detection and classification in real-world traffic conditions.

Edge vs. Cloud Implementation

A critical architectural consideration in intelligent traffic surveillance systems is the choice between edge and cloud computing for model inference and data processing. This decision directly affects performance, latency, scalability, and privacy.

In cloud computing architectures, processing is centralized on remote servers equipped with high-performance GPUs. This setup offers several advantages, including powerful computational capacity, seamless deployment of model updates, and the ability to aggregate and analyze large volumes of data for continuous system improvement. However, cloud-based approaches also introduce notable drawbacks: transmitting high-resolution video streams to the cloud increases latency, makes performance dependent on network stability, and raises privacy concerns due to the handling of sensitive vehicle and personal data over the internet.

Conversely, edge computing performs inference locally, directly on-site through embedded devices such as NVIDIA Jetson or other AI accelerators. This approach significantly reduces latency, often to below 200 milliseconds, minimizes bandwidth consumption, and enhances data privacy, as sensitive information is processed locally rather than transmitted externally. Additionally, edge systems remain functional during network outages, providing greater resilience. Nonetheless, they face challenges such as limited computational resources, higher initial hardware costs, and the complexity of maintaining and updating distributed devices in the field.

To reconcile these trade-offs, a hybrid Edge–Cloud architecture has emerged as a practical and efficient solution. In this design, immediate tasks such as vehicle detection and basic classification are handled locally at the edge, while more resource-intensive processes—like advanced analytics, retraining, and long-term data management—are offloaded to the cloud. This configuration effectively combines the low latency and privacy benefits of edge computing with the scalability and computational power of the cloud, providing a robust framework for modern, real-time traffic surveillance systems (Satyanarayanan, 2017).

Hardware Considerations

The performance of AI-driven traffic surveillance systems is heavily dependent on the underlying hardware infrastructure, as each component contributes directly to system accuracy, speed, and reliability.

At the foundation are the cameras, which serve as the system's primary sensors. Their resolution, frame rate, and low-light sensitivity determine the clarity and usability of captured footage. Advanced features such as Wide Dynamic Range (WDR) allow effective monitoring in environments with varying lighting conditions, such as bright sunlight or deep shadows, while infrared (IR) support ensures continuous, 24-hour operation even in low-visibility settings. Equally important are Graphics Processing Units (GPUs), which accelerate deep learning inference and enable real-time performance for demanding tasks like Automatic License Plate Recognition (ALPR). Without sufficient GPU capability, system latency increases, reducing the effectiveness of live monitoring and rapid violation detection.

For edge-based deployments, embedded systems such as System-on-Chip (SoC) platforms, like the NVIDIA Jetson series or Google Coral Edge TPU—offer a practical balance between power efficiency and computational capability.

These compact devices can process video streams locally, making them ideal for decentralized surveillance setups.

In addition, reliable storage and networking components are essential for maintaining data integrity. High-capacity local storage ensures temporary buffering and backup during connectivity interruptions, while stable, high-bandwidth connections allow for smooth data transmission to central servers when needed.

In the context of Albania, where cost-effectiveness is a key concern, lightweight yet capable solutions are more practical than large-scale GPU clusters. Platforms such as Jetson Nano or Jetson Xavier provide robust performance at a fraction of the cost, making them ideal for pilot projects and early-stage system deployment. These configurations balance affordability and performance, enabling sustainable development of AI-based traffic surveillance across the country.

License Plate Recognition (ALPR) and Vehicle Attribute Recognition (VAMR)

Detection: Real time ALPR Systems Automatic License Plate Recognition (ALPR) involves the detection and identification of vehicle license plates within images or video streams.

The initial stage, License Plate Detection (LPD), aims to accurately locate the region containing a license plate. The overall performance of an ALPR system heavily depends on the precision of this detection. For traffic monitoring applications, real-time detection is essential, requiring a processing speed (FPS – Frames Per Second) that can match the video feed.

Real-world traffic scenarios present multiple challenges for plate detection. Variations in lighting, from bright sunlight to shadows or low-light conditions, affect visibility. Weather conditions such as rain, fog, or snow may obscure plates. Camera angles and distances result in varying perspectives and plate sizes, while high vehicle speeds can introduce motion blur. Partial occlusions caused by trailers, other vehicles, or dirt, as well as background textures resembling plates (e.g., advertisements or signs), can generate false positives.

Traditional LPD methods relied on handcrafted features such as vertical and horizontal edges, aspect ratios, and plate colors, using tools like Sobel and Canny edge detectors or Hough line transforms. However, modern real-time ALPR systems predominantly employ single-stage deep learning models, such as YOLO, which predict bounding boxes and classes in a single pass. Recent versions of YOLO (v5, v7, v8) offer device-specific variants and advanced training techniques. Models trained on local (Albanian) and international license plate datasets can achieve high-speed, accurate detection.

FIGURE 2: Example of integrated driver detection
and license plate recognition in a real traffic environment.

OCR for different international plate styles

OCR converts the cropped license plate image into an alphanumeric string representing the vehicle's registration number. Plate formats vary considerably across regions. Some have a single row of characters, while others feature multiple rows. Fonts differ, with stylized or region-specific designs, often leading to confusion between visually similar characters ('O' vs '0', 'I' vs '1', 'B' vs '8').

Character sets also vary: some plates use only uppercase Latin letters and Arabic numerals, whereas others include lowercase letters, special characters, or symbols from non-Latin scripts such as Cyrillic, Arabic, or Chinese. Background and character colors differ (e.g., black on yellow), which can assist in localization or recognition. Plate materials range from retroreflective to non-reflective, and security features such as holograms or watermarks may affect OCR accuracy.

Robust OCR models must handle these variations, combining image preprocessing, character segmentation, and classification techniques. Deep learning approaches, particularly convolutional neural networks (CNNs) and transformer-based models, have shown superior performance in handling fonts, layout, and color diversity. Additionally, integrating contextual rules, such as expected plate formats for a given country, enhances accuracy and reduces errors caused by ambiguous characters or low-quality images.

ALPR in the Albanian Context and OCR Using Deep Learning

In the Albanian context, ALPR systems must recognize both current and historical Albanian license plate formats, as well as a wide range of international plates from neighboring countries and beyond.

After license plate detection, Optical Character Recognition (OCR) typically follows a two-step process. First, character segmentation separates each character into individual regions. Next, each character is classified using techniques such as neural networks, support vector machines (SVMs), or template matching.

Traditional approaches are sensitive to errors caused by poor image quality, complex backgrounds, or closely spaced characters.

Modern deep learning models, particularly Convolutional Recurrent Neural Networks (CRNNs), provide effective solutions. CRNNs initially apply convolutional layers to extract visual features from the plate image. These features are then processed through recurrent layers, such as LSTMs, which capture sequential dependencies among characters. A specialized CTC (Connectionist Temporal Classification) loss function enables training without precise character segmentation, making CRNNs suitable for plates with varying lengths and layouts.

Training a CRNN for international plate recognition requires a large, diverse dataset encompassing plates from multiple countries and different formats. Data augmentation, including changes in image angles, lighting conditions, or artificial noise, enhances model robustness. Additionally, transfer learning can be applied: a model pretrained on a large, general dataset (e.g., plates from other regions) is fine-tuned on a specific dataset containing Albanian and other relevant plates.

Although most models are primarily trained for Latin letters and Arabic numerals, their architecture is flexible, allowing adaptation to recognize characters from other alphabets if needed. This approach ensures accurate recognition across a wide variety of plate styles, supporting robust ALPR systems in both national and international traffic monitoring contexts.

Handling Non-Standard and Multilingual Plates

ALPR systems often face challenges beyond standard plate formats. These include vanity plates with unusual characters or spacing, damaged or faded plates, and dirty or obstructed plates that reduce visibility. Other issues arise from intentionally altered, forged, or homemade plates, as well as special types like diplomatic or temporary plates with unique designs. Handling such variability requires adaptive recognition models and preprocessing techniques to maintain accuracy in diverse real-world conditions.

Effective handling of these plates requires a robust system capable of accurate detection and OCR. Deep learning models can learn from diverse examples; training datasets that include damaged or dirty plates improve system resilience. In cases of severe damage or manipulation, partial recognition may be possible, or the plate may be flagged as "unreadable" or "suspicious," necessitating manual review or cross-referencing with vehicle data from VAMR. Some systems incorporate modules to detect signs of tampering.

Multilingual license plates add complexity to ALPR systems, as they may contain characters from different alphabets. While most European plates use the Latin script, some vehicles from the Balkans or Eastern Europe feature Cyrillic characters, requiring broader recognition capabilities. OCR models like CRNNs can handle such cases if three conditions are met: the character set includes all relevant symbols, there is adequate training data for each alphabet, and the system effectively manages visually similar characters across scripts. A "universal" OCR model may cover all expected characters, or specialized OCR models can be applied for specific alphabets based on preliminary country identification. Modern systems use Unicode to uniquely represent all characters. This study focuses primarily on Latin letters and Arabic numerals while acknowledging broader multilingual challenges.

Environmental and Operational Challenges for ALPR

The performance of ALPR systems is heavily influenced by factors that degrade image quality or obscure license plates. Image noise is a common issue, arising from camera sensors (especially in low-light conditions), video compression, or adverse weather like rain and fog. High noise levels complicate both plate localization and character recognition. Techniques such as Gaussian or median filtering help reduce noise while preserving essential details, and training models on noisy images improves robustness.

Visual obstructions (occlusion) further challenge accuracy, as plates are often partially hidden. Common sources include vehicle parts (e.g., trailer hitches, misaligned decorative frames), external objects (nearby vehicles, pedestrians, bicycles, vegetation), and dirt or debris (mud, snow, leaves). Lighting conditions, such as glare from sunlight or deep shadows, can also distort visibility, making reliable recognition difficult. Addressing these challenges requires both preprocessing enhancements and robust, well-trained detection and OCR models.

FIGURE 3: Comparison between short-exposure and long-exposure frames for improved nighttime license plate detection.

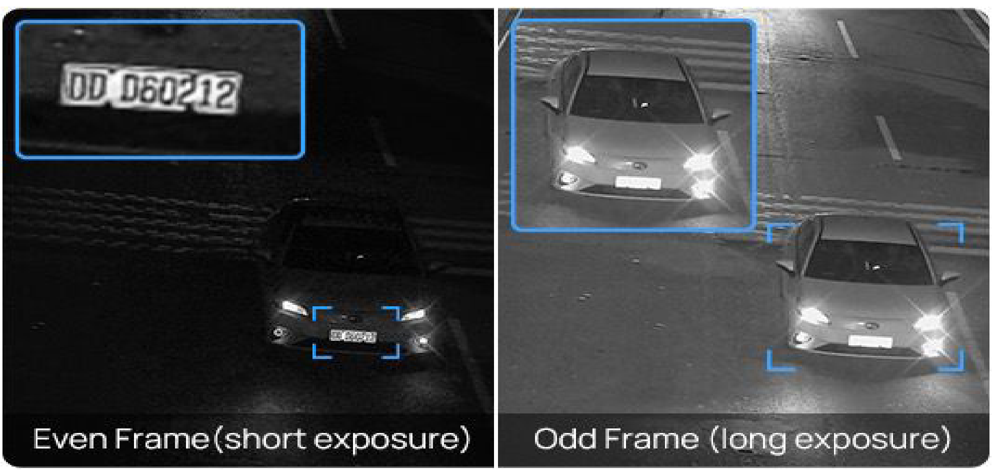

System capability to handle occlusion depends on the extent of coverage and algorithm sophistication. Modern detectors like YOLO can often locate a partially obstructed plate if distinctive features remain visible. When critical characters are occluded, full recognition may not be possible; however, some systems can perform partial reading or use contextual information to infer missing characters.

Nighttime conditions pose significant challenges for ALPR systems. Low ambient light reduces contrast and amplifies noise, while vehicle headlights and reflective plate surfaces can cause glare or overexposure. Motion blur from longer exposures further complicates character recognition.

To address these issues, infrared (IR) imaging or adaptive exposure techniques are employed. IR illumination enhances the contrast between characters and the plate background, improving detection and OCR accuracy. Successful implementation requires careful calibration of IR intensity and shutter synchronization. Additionally, including low-light and night-time images in training datasets helps models generalize better to these challenging scenarios.

FIGURE 4: License plate captures from stationary and moving vehicles using ALPR/ANPR technologies in real traffic environments.

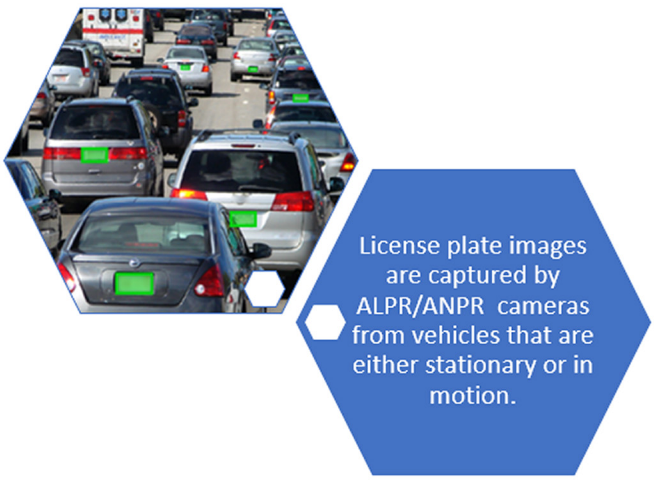

Modern traffic surveillance combines multiple AI capabilities to enhance vehicle monitoring and identification. ALPR solutions like Open ALPR, Sighthound, Plate Recognizer, and DL-based YOLO+CRNN libraries vary in accuracy, speed, coverage, and ease of integration, providing options from open-source flexibility to high-accuracy commercial systems. When plates are unreadable, Vehicle Re-Identification (Re-ID) complements ALPR by matching vehicles across cameras using CNN-based feature embeddings, handling challenges like changing viewpoints, lighting, and occlusions. Color recognition adds another layer of verification, often using HSV/HSL color spaces, SVMs, or CNNs, with a limited set of color classes to ensure consistency under varying illumination. For deeper verification, Vehicle Make and Model Recognition (VMMR) classifies specific makes and models, overcoming intra-class variation and inter-class similarity through fine-tuned pre-trained CNNs enhanced with part localization, attention mechanisms, and multi-task learning, enabling simultaneous prediction of make, model, and color. Modern traffic surveillance increasingly leverages synthetic data to enhance AI performance. Tools like Unreal Engine, Unity, CARLA, and NVIDIA DRIVE Sim generate large, fully annotated datasets with controlled variations, while techniques like domain randomization and adaptation bridge the gap between synthetic and real images.

For real-time monitoring and event response, integrated ALPR/VAMR systems connect with databases from law enforcement, vehicle registration authorities, courts, and insurance companies. This enables immediate detection of unregistered or uninsured vehicles, cross-checking against Interpol lists, and linking vehicles to their violation history.

A robust system relies on secure and scalable infrastructure, including standardized RESTful APIs, encrypted communications (WebSocket, TLS 1.3, HTTPS), OAuth2 with MFA, and full audit logs. Cloud-native or hybrid platforms allow dynamic scaling during peak traffic hours. Advanced ALPR models, such as YOLOv5 combined with CRNN, trained on Albanian-specific datasets, achieve recognition accuracy above 95%, even under low-light conditions, ensuring reliable real-time performance.

Best practices worldwide

Successful national implementations exist in South Korea, Estonia, and the UK. London's Smart Surveillance reduced vehicle-related crime by 30% .Estonia's X-Road centralized violation management, cutting penalty issuance from 48 hours to five minutes.

FIGURE 5: Improvement in plate-recognition accuracy using a 60 fps @1080p camera compared to a regular imaging sensor.

For Albania, developing an effective AI-based traffic surveillance system requires both regulatory support and technological planning. Establishing legislation that allows ALPR-generated evidence to be admissible in court is a critical first step, ensuring that automated detections have legal validity. Simultaneously, creating an integrated road safety platform that connects the DPSHTRR, police, courts, and insurance companies would enable real-time data sharing and coordinated responses to traffic violations.

From a technological perspective, edge computing should be leveraged to enable faster, localized vehicle detection and reduce latency, particularly in high-traffic areas. Where authentic Albanian datasets are limited, the use of synthetic data can help train AI models, providing diverse and annotated images to bootstrap system performance. Together, these measures would lay a strong foundation for a robust, accurate, and legally supported traffic surveillance infrastructure in Albania.

Real-World ALPR/VAMR Pipeline Architecture

Designing an effective automated traffic surveillance system involves creating a pipeline architecture that integrates multiple components into a coherent workflow. The process begins with video stream acquisition, capturing footage from multiple cameras via standard protocols like RTSP. These streams are then decoded into individual frames and preprocessed—such as resizing or normalization—to prepare the data for AI models.

Next, vehicle and license plate detection localizes relevant objects in each frame, typically using models like YOLO. Detected vehicles are assigned unique IDs and followed across frames through multi-object tracking algorithms such as DeepSORT, ensuring temporal consistency. From these detections, Regions of Interest (ROIs) are cropped for specialized analysis.

Automatic License Plate Recognition (ALPR) processes cropped plate images using OCR models (e.g., CRNN with CTC loss) to extract alphanumeric strings, while Vehicle Attribute Recognition (VAMR) analyzes cropped vehicle images to determine attributes such as color, make, and model using CNN-based architectures. The outputs from ALPR and VAMR are then aggregated and verified, optionally cross-checked against databases for consistency.

For enhanced functionality, the pipeline can include violation detection, identifying stolen or wanted vehicles, speeding (via radar/LiDAR integration), traffic signal infractions, or illegal parking. The system then generates evidence packages containing images, videos, plate information, vehicle attributes, timestamps, and location data, sending real-time alerts when necessary. Finally, all data and metadata are securely stored in an organized fashion, supporting auditing, analysis, and long-term monitoring.

This end-to-end pipeline ensures accuracy, efficiency, and reliability, providing a robust foundation for modern, AI-driven traffic surveillance systems.

FIGURE 6: Example of input video frame (top) and the corresponding detection and tracking output produced by the implemented ALPR pipeline (bottom).

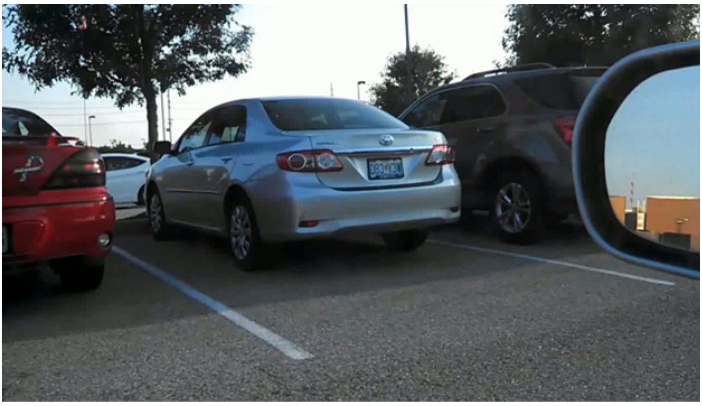

Photo from the program executed by the code

Cost Analysis for Automated Traffic Monitoring Systems

Analyzing the costs and operational considerations of automated traffic surveillance systems requires evaluating hardware, maintenance, energy consumption, and integration capabilities.

Edge computing offers low-latency, real-time processing through devices like Google Coral Edge TPU or NVIDIA Jetson Orin NX, but deploying multiple monitoring points involves significant initial investment. Operational costs include energy consumption, regular maintenance, software updates, and hardware replacement. By processing data locally, edge devices reduce bandwidth usage and network congestion, sending only summarized outputs such as violation events, timestamps, and locations. Environmentally, edge computing can leverage renewable energy sources like solar power—especially practical in sunny regions such as Albania—reducing reliance on energy-intensive cloud data centers.

Case studies demonstrate diverse implementations: Dubai uses AI-based ALPR for speeding and traffic safety monitoring; Singapore integrates ALPR into intelligent transport systems for traffic optimization and electronic road pricing; London applies ALPR extensively for congestion and low-emission zones, though high maintenance costs and privacy concerns remain challenges.

Integration with national databases and law enforcement systems is essential for maximizing system effectiveness. Linking ALPR outputs with the Vehicle Registration Database (VRD), such as Albania's DPSHTRR, allows verification of ownership, insurance status, technical inspections, and vehicle attributes, facilitating automated fine issuance and detection of cloned plates. Connecting to police databases enables alerts for stolen or wanted vehicles, historical violations, and links to wanted persons, supporting rapid law enforcement intervention.

Key challenges in integration include the availability and standardization of APIs, data security and access control, performance and scalability, data quality, and establishing a legal framework for data sharing. Addressing these factors ensures that ALPR/VAMR systems operate efficiently, securely, and reliably while supporting law enforcement and traffic management objectives.

Data Storage and Security

Effective automated traffic monitoring systems require comprehensive data management and security strategies. Data retention policies should define how long information is stored, ensuring that unnecessary data is securely deleted or anonymized in compliance with legal requirements. Access control mechanisms—authentication, authorization, and audit log—restrict data access to authorized personnel and track activity. Encryption protects sensitive data both in transit and at rest, while physical and virtual infrastructure security safeguards hardware and cloud environments. Techniques like anonymization and pseudonymization reduce privacy risks when data is used for analysis or model training. Predefined incident response plans ensure rapid mitigation of breaches, and ongoing personnel training reinforces secure handling practices.

Automated Violation Decision Support Systems (AVDSS)

AVDSS use advanced algorithms, such as Random Forests, XGBoost, and neural networks—to detect, classify, and respond to traffic violations, often achieving over 95% accuracy. Components include violation detection, categorization by severity, and automated decision-making (issuing fines, alerting authorities, storing evidence). Visualization tools like GIS maps and temporal graphs support urban planning, while machine learning techniques (supervised and unsupervised) predict risky behaviors and detect patterns. CNNs can visually identify dangerous driving, and continuous learning improves system performance over time. Key risks involve algorithmic bias, misclassifications, and overreliance on automation, highlighting the need for transparency, auditability, and human oversight.

Conclusion

AI-powered AVDSS enhances road safety, operational efficiency, and urban planning by accurately detecting and analyzing traffic violations. While challenges such as bias, complex data, and ethical concerns persist, continuous learning and human supervision ensure fairness and reliability. These systems not only enforce traffic laws but also contribute to smarter, safer cities.

Key Findings and Recommendations

- Transformative Potential of AI: Deep learning and computer vision models (YOLO, CRNN, CNNs for VAMR) provide high accuracy and speed for automated traffic monitoring.
- Integrated ALPR+VAMR Architecture: Combining license plate and vehicle attribute recognition improves reliability, especially when plates are unreadable.
- Real-World Challenges: Lighting, weather, occlusions, high speeds, and diverse plate formats require robust datasets, augmentation, and advanced models.
- Critical Role of Data: High-quality, diverse, locally annotated datasets are essential; their scarcity is a major barrier.
- Edge vs. Cloud Trade-offs: Edge computing reduces latency and bandwidth usage, while cloud systems enable centralized management and deeper analysis.
- Ethical and Legal Considerations: Privacy, accountability, and transparency are crucial; compliance with legal frameworks (e.g., GDPR) and ethical AI principles is mandatory.

Best Practices

- Modular and Integrated Approach: Building the system as interconnected modules (detection, tracking, ALPR, VAMR) while maintaining a coherent pipeline for optimal performance.

- Advanced Models and Fine-Tuning: Leveraging pretrained models (YOLO, ResNet, CRNN) and fine-tuning them on task-specific, local datasets.
- Focus on Robustness: Training models to withstand variations in lighting, weather, occlusion, and vehicle/plate types using data augmentation and diverse datasets.
- Continuous Evaluation: Using appropriate performance metrics, benchmark datasets, and real-world testing to identify and address weaknesses.
- Hybrid Edge-Cloud Architecture: Edge processing for time-critical tasks and cloud for deep analysis and centralized storage.
- Privacy and Ethics by Design: Integrating ethical and privacy considerations early in system design.
- Interinstitutional Collaboration: Coordinating among police, transport authorities, and data protection agencies.
- Stakeholder Engagement: Involving the public, civil society, and domain experts to ensure acceptability and trust.

Recommendations for Albania

- Start with Limited Pilot Projects: Test technology in selected high-risk areas to collect context-specific data and assess effectiveness.
- Invest in Local Data Collection and Annotation: Develop large, diverse datasets of traffic images from Albanian roads.
- Adopt Open and Flexible Architectures: Use open-source, well-tested models that can be fine-tuned for local conditions.
- Prioritize Low Latency for Critical Applications: Use edge processing for rapid response tasks such as stolen vehicle detection.
- Develop Clear Legal and Ethical Frameworks: Update legislation to regulate AI-based traffic monitoring in compliance with GDPR and human rights standards.
- Strengthen Technical and Human Capacity: Train technical staff and law enforcement personnel for correct and ethical system use.

Future Research Directions

- Enhancing Robustness in Extreme Conditions: Develop ALPR/VAMR models resilient to snow, low light, or severe occlusions, possibly using multimodal sensors (visual, radar, Lidar).
- Few-Shot and Self-Supervised Learning: Enable recognition of new plate formats or vehicle models with minimal labeled data.
- Continual Domain Adaptation: Create models that adapt dynamically to environmental changes without full retraining.

- Explainable AI (XAI) for Traffic Monitoring: Improve model transparency and interpretability to increase trust and facilitate error analysis.
- Advanced License Plate Forgery Detection: Detect subtle modifications or anti-ALPR materials.
- Vehicle Behavior Analysis: Use AI to predict risks and detect unsafe driving patterns beyond plate and attribute recognition.
- Optimization for Resource-Constrained Edge Devices: Develop energy-efficient and low-cost deep learning solutions for edge deployment.
- Socio-Ethical Studies: Conduct empirical research on public perception, societal impacts, and context-specific regulatory frameworks.

References

Anagnostopoulos, C. N. E., Anagnostopoulos, I. E., Psoroulas, I. D., Loumos, V., & Kayafas, E. (2008). License plate recognition from still images and video sequences: A survey. *IEEE Transactions on Intelligent Transportation Systems, 9*(3), 377–391. https://doi.org/10.1109/TITS.2008.922938

Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and machine learning. MIT Press. https://fairmlbook.org/

Baek, J., Lee, B., Han, D., Yun, S., & Lee, H. (2019). Character region awareness for text detection.

Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 12275–12283). Retrieved from https://openaccess.thecvf.com/content_CVPR_2019/html/Baek_Character_Region_Awareness_for_Text_Detection_CVPR_2019_paper.html

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification.

Proceedings of the Conference on Fairness, Accountability and Transparency (pp. 77–91). Retrieved from https://dl.acm.org/doi/10.1145/3287560.3287565

Chen, L., et al. (2021). A survey on intelligent transportation systems. *IEEE Transactions on Intelligent Transportation Systems, 22*(2), 618–637. https://doi.org/10.1109/TITS.2020.2988061

European Transport Safety Council. (2022, June 15). The 2022 ETSC Road Safety Performance Index Conference. https://etsc.eu/the-2022-etsc-road-safety-performance-index-conference/

European Union Agency for Fundamental Rights. (2021). *Getting the future right: Artificial intelligence and fundamental rights*. Publications Office of the European Union. https://fra.europa.eu

Goh, P. H., et al. (2022). Intelligent transportation systems: A review of technologies and applications. *Journal of Traffic and Transportation Engineering (English Edition), 9*(1), 1–25. Retrieved from https://www.sciencedirect.com/science/article/pii/S####

Instituti i Statistikave (INSTAT). (2023). *Statistika mbi sigurinë rrugore*. Tiranë, Shqipëri. Retrieved from: https://www.instat.gov.al/media/14132/statistika-mbi-sigurine-rrugore-2023.pdf

Instituti i Statistikave (INSTAT). (n.d.). *Transporti, aksidentet dhe karakteristikat e mjetet rrugore*. Retrieved from: https://www.instat.gov.al/sq/temat/industria-tregtia-dhe-sherbimet/transporti-aksidentet-dhe-karakteristikat-e-mjetet-rrugore/

Kornblith, S., Shlens, J., & Le, Q. V. (2019). Do better ImageNet models transfer better? Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2661–2671. https://doi.org/10.1109/CVPR.2019.00277.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539.

Li, H., et al. (2019). Vehicle recognition using multi-task learning for joint make, model, and color classification. *IEEE Transactions on Intelligent Transportation Systems, 20*(3), 1019–1028. https://doi.org/10.1109/TITS.2018.2873130

Park, S. H., et al. (2020). AI-based real-time traffic violation detection system using CCTV images. *Electronics, 9*(11), 1878. https://doi.org/10.3390/electronics9111878.

Redmon, J., & Farhadi, A. (2018). *YOLOv3: An incremental improvement* [Preprint]. arXiv. https://arxiv.org/abs/1804.02767.

Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9.

Selmi, Z., Halima, M. B., & Alimi, A. M. (2017). Deep learning system for automatic license plate detection and recognition. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)* (Vol. 1, pp. 1140–1145). Retrieved from https://ieeexplore.ieee.org/document/8433936

Shi, B., Bai, X., & Yao, C. (2016). An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 39*(11), 2298–2304. https://doi.org/10.1109/TPAMI.2016.2515971

Shi, W., et al. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal, 3*(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198

Tehrani, M. H., et al. (2019). A review on vision-based license plate recognition systems. *IEEE Transactions on Intelligent Transportation Systems, 20*(8), 2881–2901. https://doi.org/10.1109/TITS.2018.2875015.

Wojke, N., Bewley, A., & Paulus, D. (2017). Simple online and realtime tracking with a deep association metric. In 2017 IEEE International Conference on Image Processing (ICIP) (pp. 3645–3649). IEEE. https://doi.org/10.1109/ICIP.2017.8296962

World Health Organization. (2023, May 15). Transport systems need to be made safe, healthy and sustainable. Retrieved from: https://www.who.int/news/item/15-05-2023-transport-systems-need-to-be-made-safe--healthy-and-sustainable

World Health Organization. (2023). *Global status report on road safety 2023*.

Retrieved from https://www.who.int/publications/i/item/9789240086517

Yang, L., Luo, P., Change Loy, C., & Tang, X. (2015). *A large-scale car dataset for fine-grained categorization and verification*. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 3973–3981). https://doi.org/10.1109/CVPR.2015.7299023

Zhang, G., Wang, Y., Wei, H., & Chen, Y. (2017). A comprehensive review of radar and LiDAR sensors for traffic enforcement and vehicle detection. *IEEE Transactions on Intelligent Transportation Systems, 18*(10), 2848–2863. https://doi.org/10.1109/TITS.2017.2680468